An Explicit Lower Bound for TSP with Distances One and Two

Lars Engebretsen

Dept. of Numerical Analysis and Computing Science
Royal Institute of Technology
SE-100 44 Stockholm, SWEDEN
E-mail: enge@nada.kth.se

January 1999

Abstract. We show that it is, for any $\epsilon > 0$, **NP**-hard to approximate the asymmetric traveling salesman problem with distances one and two within $2805/2804 - \epsilon$. For the special case where the distance function is constrained to be symmetric, we show a lower bound of $5381/5380 - \epsilon$, for any $\epsilon > 0$. While it was previously known that there exists some constant, strictly greater than one, such that it is **NP**-hard to approximate the traveling salesman problem with distances one and two within that constant, this result is a first step towards the establishment of a good bound.

In our proof, we develop a new gadget construction to reduce from systems of linear equations mod 2 with two unknowns in each equation and at most three occurrences of each variable. Compared with earlier reductions to the traveling salesman problem with distances one and two, ours reduces the number of cities to less than a tenth of what was previously necessary.

Key words. The traveling salesman problem, Approximability, Gadget reductions, Lower bounds.

1 Introduction

A common special case of the traveling salesman problem is the metric traveling salesman problem, where the distances between the cities satisfy the triangle inequality. In this paper, we study a further specialization: The traveling salesman problem with distances one and two between the cities. This problem was shown to be **NP**-complete by Karp [7]. Since this means that we have little hope of computing exact solutions, it is interesting to try to find an approximate solution, i.e., a tour with weight close to the optimum weight. Christofides [2] has constructed an elegant algorithm approximating the metric traveling salesman problem within 3/2. This algorithm also applies to the symmetric traveling salesman problem with distances one and two, but it is possible to do better; Papadimitriou and Yannakakis [8] have shown that it is possible to approximate the latter problem within 7/6.

They also show a lower bound; that there exists some constant, which is never given explicitly in the paper, such that it is **NP**-hard to approximate the problem within that constant. This hardness result extends also to the asymmetric version of the traveling salesman problem with distances one and two. As for the approximability of the asymmetric version, a 17/12-approximation algorithm has been provided by Vishwanatan [10].

Recently, there has been a renewed interest in the hardness of approximating the traveling salesman problem with distances one and two. Fernandez de la Vega and Karpinski [3] and, independently, Fotakis and Spirakis [4] have shown that the hardness result of Papadimitriou and Yannakakis holds also for dense instances. We contribute to this line of research by showing an explicit lower bound on the approximability. More specifically, we construct a reduction from linear equations mod 2 with three occurrences of each variable to show that it is **NP**-hard to approximate the asymmetric traveling salesman problem with distances one and two within $2805/2804 - \epsilon$. For the symmetric version of the problem we show a lower bound of $5381/5380 - \epsilon$.

2 Preliminaries

Definition 1. Let P be an **NP** minimization problem. For an instance x of P let $\operatorname{opt}(x)$ be the optimal value. A solution y, with weight $w(x, y)$, is *c-approximate* if it is feasible and $w(x, y)/\operatorname{opt}(x) \leq c$.

Definition 2. A *c-approximation algorithm* for an **NP** optimization problem P is a polynomial time algorithm that for any instance x and any input y outputs a c-approximate solution.

We use the wording *to approximate within c* as a synonym for *to compute a c-approximate solution.*

Definition 3. *E2-Lin mod* 2 is the following maximization problem: Given a system of linear equations mod 2 with exactly two variables in each equation, maximize the number of satisfied equations. *E2-Lin(3) mod* 2 is the special case of E2-Lin mod 2 where there are exactly three occurrences of each variable.

Definition 4. *Asymmetric (1,2)-TSP* is the special case of the traveling salesman problem where the off-diagonal entries in the distance matrix are either one or two. *Asymmetric* $\triangle$*-TSP* is the special case of the traveling salesman problem where the entries in the distance matrix obey the triangle inequality.

Definition 5. *(1,2)-TSP* and $\triangle$*-TSP*, respectively, are the special cases of *Asymmetric (1,2)-TSP* and *Asymmetric* $\triangle$*-TSP*, respectively, where the distance matrix is symmetric.

We note in passing, that since (1,2)-TSP is a special case of △-TSP a lower bound on the approximability of (1,2)-TSP is also a lower bound on the approximability of △-TSP. The same holds for the asymmetric versions of the problems.

To describe a (1,2)-TSP instance, it is enough to specify the edges of weight one. We do this by constructing a graph G, and then let the (1,2)-TSP instance have the nodes of G as cities. The distance between two cities u and v is defined to be one if (u, v) is an edge in G and two otherwise. To compute the weight of a tour, it is enough to study the parts of the tour traversing edges of G. In the asymmetric case, G is a directed graph.

Definition 6. We call a node where the tour leaves or enters G an *endpoint*. A city with the property that the tour both enters and leaves G in that particular city is called a *double endpoint*, and counts as two endpoints.

If c is the number of cities and $2e$ is the total number of endpoints, the weight of the tour is $c + e$, since every edge of weight two corresponds to two endpoints. When we analyze our reduction, we study an arbitrary tour restricted to certain subgraphs of G. Generally, such a restriction consists of several disjoint paths. To shorten the notation, we call these paths *partial tours*.

3 The symmetric case

To obtain our hardness result we reduce from E2-Lin(3) mod 2. Previous reductions from integer programming [6] and satisfiability [8] to (1,2)-TSP make heavy use of the so called *xor gadget*. This gadget is used both to link variable gadgets with equation gadgets and to obtain a consistent assignment to the variables in the original instance. The xor gadget contains twelve nodes, which means that a gadget containing some twenty xor gadgets for each variable — which is actually the case in the previously known reductions — produces a very poor lower bound. To obtain a reasonable inapproximability result, we modify the previously used xor gadget to construct an *equation gadget*. Each occurrence of a variable corresponds to a node in an equation gadget. Since each variable occurs three times, there are three nodes corresponding to each variable. These nodes are linked together in a *variable cluster*. The idea behind this is that the extra edges in the cluster should force the nodes to represent the same value for all three occurrences of the variable. This construction contains 24 nodes for each variable, which is a vast improvement compared to earlier constructions.

We give our construction in greater detail below. In a sequel of lemmas, we show that an optimal tour can be assumed to have a certain structure. We do this by showing that we can transform, by local transformations which do not increase the length of the tour, any tour into a tour with the sought

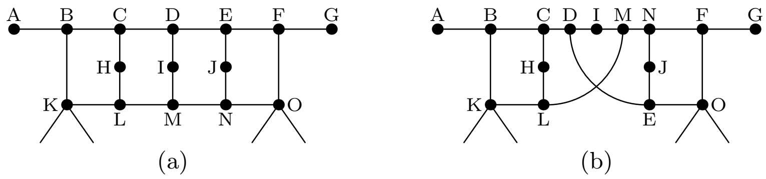

Figure 1. The equation gadget is connected to other gadgets through the vertices A and G and through the edges shown above from the vertices K and O. Gadget (a) corresponds to an equation of the form $x + y = 1$ and gadget (b) to an equation of the form $x + y = 0$.

structure. This new tour, obtained after the local transformations, can then be used to construct an assignment to the variables in the original E2-Lin(3) mod 2 instance. Our main result follows from a recent hardness result of Berman and Karpinski [1] together with a correspondence between the length of the tour in the (1,2)-TSP instance and the number of unsatisfied equations in the E2-Lin(3) mod 2 instance.

3.1 The equation gadget

The equation gadget is shown in Fig. 1. It is connected to other gadgets in four places. The vertices A and G coincide with similar vertices at other gadgets to form a long chain. Thus, these vertices have degree two. The edges from the vertices K and O join the equation gadget with other equation gadgets. We study this closely in Sec. 3.2. No other vertex in the gadget is joined with vertices not belonging to the gadget.

Definition 7. We from now on call the vertices K and O in Fig. 1 the *lambda-vertices* of the gadget and the boundary edges connected to these vertices *lambda-edges*. For short, we often refer to the pair of lambda-edges linked to a particular lambda-vertex as a *lambda*.

Definition 8. We say that a lambda is *traversed* if both lambda-edges are traversed by the tour, *untraversed* if none of the lambda-edges are traversed, and *semitraversed* otherwise.

In the following lemmas, we show that an optimal tour can be assumed to traverse the equation gadget in a very special way.

Lemma 9. *Suppose that we have a tour traversing an equation gadget of the type shown in Fig. 1a in such a way that there are no semitraversed lambdas in it. If there is exactly one traversed lambda, it is possible to modify this tour, without increasing its length and without changing the tour on the lambdas, in such a way that there are no endpoints in the gadget. Otherwise, it is possible to construct a tour with two endpoints in the gadget and impossible to construct a tour with less than two endpoints in the gadget.*

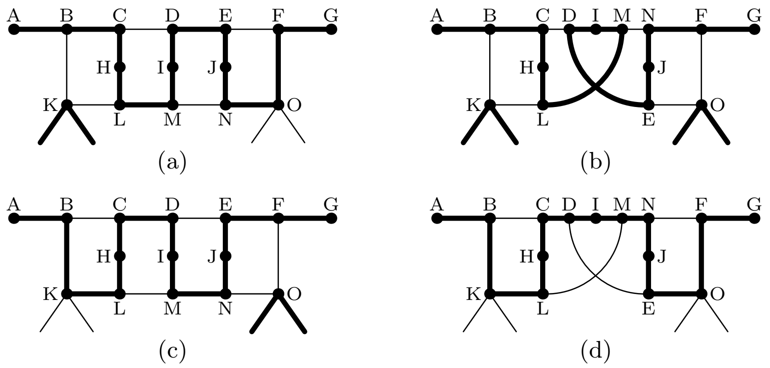

Figure 2. Given that the lambda-edges are traversed as shown above, it is possible to construct a tour through the equation gadget such that there are no endpoints in the gadget.

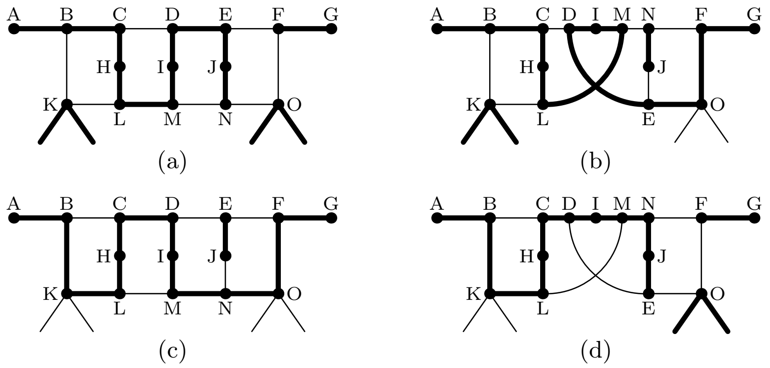

Figure 3. Given that the lambda-edges are traversed as shown above, there must be at least two endpoints in the gadget. There are in fact many tours with this property, we show only a few above.

Suppose that we have a tour traversing an equation gadget of the type shown in Fig. 1b in such a way that there are no semitraversed lambdas in it. If there are zero or two traversed lambdas, it is possible to modify this tour, without increasing its length and without changing the tour on the lambdas, in such a way that there are no endpoints in the gadget. Otherwise, it is possible to construct a tour with two endpoints in the gadget and impossible to construct a tour with less than two endpoints in the gadget.

Proof. Figures 2 and 3 show that there exist tours with the number of endpoints stated in the lemma. To complete the proof, we must show that it is impossible to construct better tours in the cases where the tour has two endpoints in the gadget. It is locally optimal to let the tour traverse the edge AB and the edge FG in Fig. 1. Thus we can assume that one partial tour enters the gadget through the vertex A and that another, or possibly the same, partial tour enters through the vertex G. Since there are

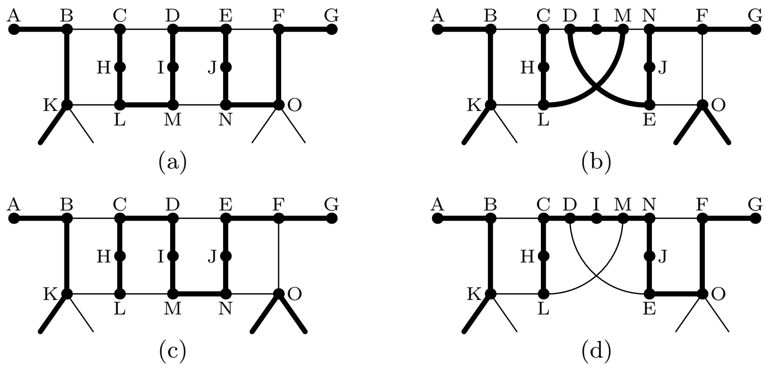

Figure 4. If one lambda is semi-traversed, there must be at least one endpoint in the gadget.

no semitraversed lambdas in the gadget, the only way, other than through the above described vertices, a partial tour can leave the gadget is through an endpoint, which in turn implies that there is an even number of endpoints in the gadget.

If there is to be no endpoint in the gadget, all of the edges CH, HL, DI, IM, EJ, and JN must be traversed by the tour. Also, the edges CD and LM cannot be traversed simultaneously, neither can DE and MN. The only way we can avoid making the vertices D and M endpoints is to traverse either the edges CD and MN or the edges DE and LM.

Let us suppose that the lambdas in the gadget are traversed as shown in Fig. 3a. By our reasoning above and the symmetry of the gadget, we can assume that the edges AB, CH, HL, LM, MI, ID, DE, EJ, JN, and FG are traversed by the tour. To avoid making the vertices C and N endpoints, the tour must traverse the edges BC and NO. But this is impossible, since the right lambda is already traversed. Thus, there is no tour with zero endpoints in the gadget, which implies that is impossible to construct a tour with less than two endpoints in the gadget. With a similar argument, we conclude that the same holds for the other cases shown in Fig. 3. □

Lemma 10. *Suppose that we have a tour traversing an equation gadget in such a way that there is exactly one semitraversed lambda in it. Then it is possible to modify this tour, without increasing its length and without changing the tour on the lambdas, in such a way that there is one endpoint in the gadget, and it is impossible to construct a tour with less than one endpoint in the gadget.*

Proof. From Fig. 4 we see that we can always construct tours such that there is one endpoint in the gadget. We now show that it is impossible to construct a tour with fewer endpoints. As in the proof of Lemma 9, we can assume that one partial tour enters the gadget at A and that another, or the

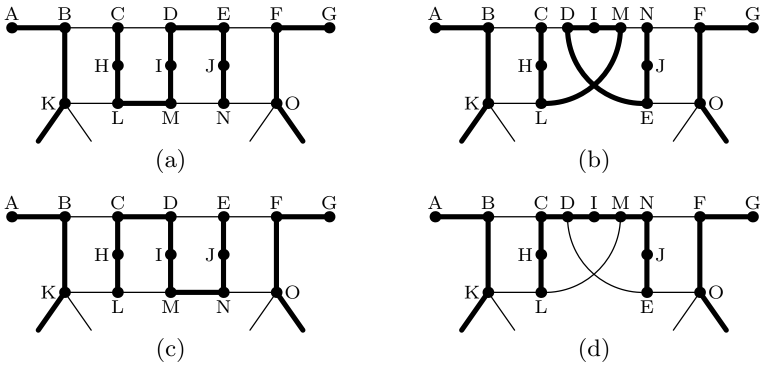

Figure 5. If both lambdas in a gadget are semi-traversed, there must be at least two endpoints in the gadget.

same, enters at G. Since there is one semitraversed lambda in the gadget, one partial tour enters the gadget at that lambda, which implies that there must be an odd number of endpoints in the gadget. □

Lemma 11. *Suppose that we have a tour traversing an equation gadget in such a way that there are two semitraversed lambdas in it. Then it is possible to modify this tour, without increasing its length and without changing the tour on the lambdas, in such a way that there are two endpoints in the gadget, and it is impossible to construct a tour with less than two endpoints in the gadget.*

Proof. From Fig. 5 we see that we can always construct tours such that there are two endpoints in the gadget. In order to prove the last part of the lemma we must argue that it is impossible to traverse the gadget in such a way that there are no endpoints in it. By an argument similar to that present in the proofs of Lemmas 9 and 10, a partial tour will enter (or leave) the gadget at four places, which implies that there is an even number of endpoints in the gadget. If there is to be no endpoints in the graph, there must be two partial tours in the gadget. Since the tours cannot cross each other and the gadget is planar we have two possible cases.

The first case is that the partial tour entering the gadget at A leaves it at G (for a gadget of the type shown in Fig. 1a) or at O (for a gadget of the type shown in Fig. 1b) and the partial tour entering at K leaves at O (for a gadget of the type shown in Fig. 1a) or at G (for a gadget of the type shown in Fig. 1b). These two partial tours cannot, however, traverse all of the edges CH, HL, DI, IM, EJ, and JN without crossing or touching each other. As noted in the proof of Lemma 9, all three abovementioned edges must be traversed for the gadget to contain no endpoints. Thus, we can rule this case out.

The second case is that the partial tour entering the gadget at A leaves it at K and the partial tour entering at G leaves at O. Since these two partial tours cannot traverse all of the edges CH, HL, DI, IM, EJ, and JN without crossing or touching each other, we conclude that at least two endpoints must occur within the equation gadget. □

Lemma 12. *It is always possible to change a semitraversed lambda to either a traversed or an untraversed lambda without increasing the number of endpoints in the tour.*

Proof. First suppose that only one of the lambdas in the equation gadget is semitraversed. By Lemma 10 we can assume that the gadgets are traversed according to Fig. 4. Let us study the tour shown in Fig. 4a. By replacing it with the tour shown in Fig. 2a, we remove one endpoint from the equation gadget, but we may in that process introduce one endpoint somewhere else in the graph. In proof, let λ be the left lambda-vertex in Fig. 4a and v be the vertex adjacent to λ through the untraversed lambda edge. If v is an endpoint, we simply let the partial tour ending at v continue to λ, thereby saving one endpoint. If v is not an endpoint, we have to reroute the tour at v to λ. This introduces an endpoint at a neighbor of v, but that endpoint is set off against the endpoint removed from the equation gadget. To sum up, we have shown that it is possible to convert the tour in Fig. 4a to the one in Fig. 2a without increasing the total number of endpoints in the graph. In a similar way, we can convert the tour in Fig. 4b to the one in Fig 2b, the tour in Fig. 4c to the one in Fig 2c, and the tour in Fig. 4d to the one in Fig 2d, respectively.

Finally, suppose that both lambdas are semitraversed. By Lemma 11 we can assume that the gadgets are traversed according to Fig. 5. By the method described in the previous paragraph we can convert the tour in Fig. 5a to the one in Fig 4a, the tour in Fig. 5b to the one in Fig 4b, the tour in Fig. 5c to the one in Fig 4c, and the tour in Fig. 5d to the one in Fig 4d, respectively. □

3.2 The variable cluster

The variable cluster is shown in Fig. 6. The vertices A and B coincide with similar vertices at other gadgets to form a long chain, as described in Sec. 3.3. Suppose that the variable cluster corresponds to some variable x. Then the upper three vertices in the cluster are lambda-vertices in the equation gadgets corresponding to equations where x occurs. The remaining two vertices in the cluster are not joined with vertices outside the cluster.

Lemma 13. *Suppose that we have a tour traversing a cluster in such a way that there are some semitraversed lambdas in it. Then, it is possible to modify the tour, without making it longer, in such way that there are no semitraversed lambdas in the cluster.*

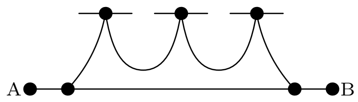

Figure 6. There is one lambda-vertex for each occurrence of each variable. The three lambda-vertices corresponding to one variable in the system of linear equations are joined together in a variable cluster. The three uppermost vertices in the figure are the lambda-vertices.

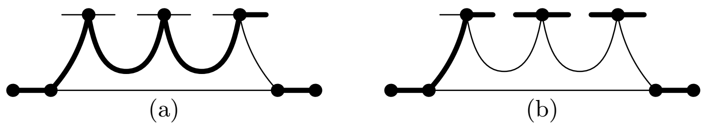

Figure 7. If the first or the last lambda in a variable cluster is semitraversed, the cluster is traversed as shown above.

Proof. Suppose that there is one semitraversed lambda. If the semitraversed lambda is the middle lambda of the variable cluster it can, by Lemma 12, be transformed into either a traversed or an untraversed lambda. This moves the semitraversed lambda to the end of the cluster. Then the cluster looks as in Fig. 7. By moving the endpoint in the variable cluster to the equation gadget corresponding to the semitraversed lambda, we can make the last semitraversed lambda traversed or untraversed without changing the number of endpoints.

Suppose now that there are two semitraversed lambdas. By Lemma 12, they can be transformed into either a traversed or an untraversed lambda without changing the number of endpoints in the tour. This implies that we can transform the tour in such a way that there is only one semitraversed lambda without changing the number of endpoints in the tour. Then we can use the method from the above paragraph to transform that tour in such a way that there are no semitraversed lambdas.

Finally, suppose that all three lambdas are semitraversed. Then the variable cluster would be traversed as in Fig. 8. By Lemma 12, the tour can be transformed in such a way that the two outer lambdas in the variable cluster are either traversed or untraversed without changing the weight of the tour. If the center lambda is not semitraversed after the transformation, the proof is complete. Otherwise we can apply the first paragraph of this proof. □

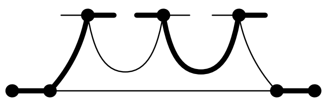

Figure 8. If all lambdas in a variable cluster are semitraversed, the cluster is traversed as shown above.

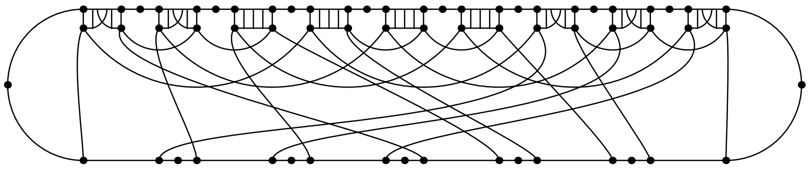

Figure 9. All equation gadgets and all variable clusters are linked together in a circular chain as shown schematically above. The equation gadgets are at the top of the figure and the variable clusters at the bottom. The precise order of the gadgets is not important. For clarity, we have omitted nine vertices from each equation gadget.

3.3 The entire (1,2)-TSP instance

To produce the (1,2)-TSP instance, all equation gadgets are linked together in series, followed by all variable clusters. The first equation gadget is also linked to the last variable cluster. The construction is shown schematically in Fig. 9. The precise order of the individual gadgets in the chain is not important.

Our aim in the analysis of our construction is to show that we can construct, from a tour containing e edges of weight two, an assignment to the variables such that at most e equations are not satisfied. When we combine this with the recent hardness results for E2-Lin(3) mod 2 [1], we obtain Theorem 15, our main result.

Lemma 14. *Given a tour with* $2e$ *endpoints, we can construct an assignment leaving at most* e *equations unsatisfied.*

Proof. Given a tour, we can by Lemmas 9–13 construct a new tour, without increasing its length, such that for each variable cluster either all or no lambda-edges are traversed. Then we can construct an assignment as follows: If the lambda-edges in a cluster are traversed by the tour, the corresponding variable is assigned the value one; otherwise it is assigned zero. By Lemma 9, this assignment has the property that there are two endpoints in the equation gadgets corresponding to unsatisfied equations. Thus, the assignment leaves at most e equations unsatisfied if there are $2e$ endpoints. □

Theorem 15. *For every* $\epsilon_1 > 0$ *and* $\epsilon_2 > 0$ *it is* ***NP****-hard to decide whether an instance of the traveling salesman problem with distances one and two with* $5376n$ *nodes has an optimum tour with length above* $(5381 - \epsilon_1)n$ *or below* $(5380 + \epsilon_2)n$.

Corollary 16. *It is, for any* $\epsilon > 0$, ***NP****-hard to approximate (1,2)-TSP within* $5381/5380 - \epsilon$.

Proof of Theorem 15. The result of Berman and Karpinski [1] states that it is **NP**-hard to determine if an instance of E2-Lin(3) mod 2 with $336n$ equations has its optimum above $(332-\epsilon_2)n$ or below $(331+\epsilon_1)n$. If we construct

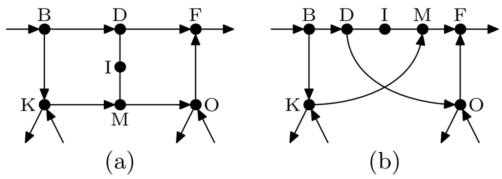

Figure 10. The asymmetric equation gadget is connected to other gadgets through the edges shown above. Gadget (a) corresponds to an equation of the form $x + y = 1$ and gadget (b) to an equation of the form $x + y = 0$.

from an instance of E2-Lin(3) mod 2 an instance of (1,2)-TSP as described above, the graph in the (1,2)-TSP instance contains $48n$ nodes given that the E2-Lin(3) mod 2 instance contains $2n$ variables and $3n$ equations. Thus, Lemma 14 and the above hardness result together imply that it is **NP**-hard to decide whether an instance of (1,2)-TSP with $5376n$ nodes has an optimum tour with length above $(5381 - \epsilon_1)n$ or below $(5380 + \epsilon_2)n$. □

4 The asymmetric case

When we show our lower bound in Sec. 3 we construct, for an arbitrary instance of E2-Lin(3) mod 2, a graph G corresponding to it. This graph has a very special structure and the purpose of our construction is to force the tour to take a certain way through graph. We then use the tour to construct an assignment to the variables, and use our knowledge of the relationship between the length of the tour and the number of unsatisfied equations to obtain a lower bound.

It seems that it would be easier to force the tour to take a certain way if we were to allow an asymmetric distance function. If we turn our attention to Asymmetric (1,2)-TSP, this means that the graph G becomes directed, i.e., it is only favorable to traverse the edges in one specific direction. It turns out that it is possible to use this to obtain a better lower bound for the directed case. The construction is very similar to the construction in Sec 3, but we repeat the details for the sake of completeness.

4.1 The equation gadget

The asymmetric equation gadget is shown in Fig. 10. The edge to B in one equation gadget coincides with the edge from F in another. Thus, the gadgets are connected in a long chain. This is described in greater detail in Sec. 4.3. The edges to and from the vertices K and O join the equation gadget with other equation gadgets as described in Sec. 4.2.

Lemma 17. *Suppose that we have a tour traversing an equation gadget of the type shown in Fig. 10a in such a way that there are no semitraversed lambdas in it. If there is exactly one traversed lambda, it is possible to*

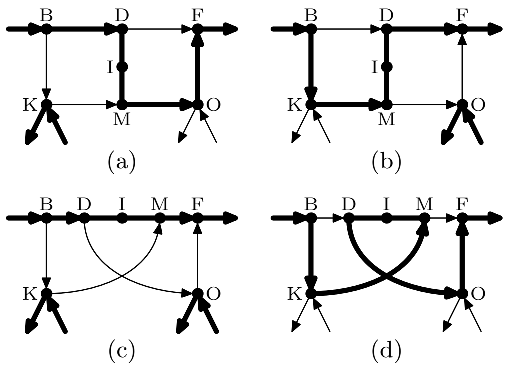

Figure 11. Given that the lambda-edges are traversed as shown above, it is possible to construct a tour through the equation gadget such that there are no endpoints in the gadget.

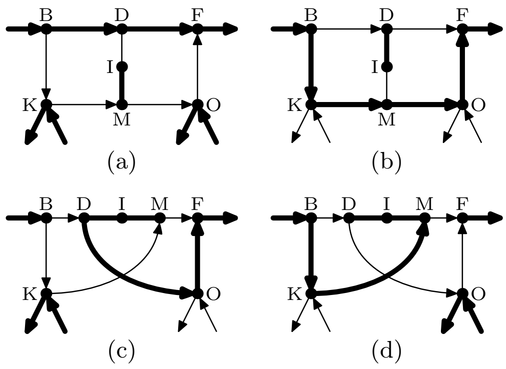

Figure 12. Given that the lambda-edges are traversed as shown above, there must be at least two endpoints in the gadget. There are in fact many tours with this property, we show only a few above.

modify this tour, without increasing its length and without changing the tour on the lambdas, in such a way that there are no endpoints in the gadget. Otherwise, it is possible to construct a tour with two endpoints in the gadget and impossible to construct a tour with less than two endpoints in the gadget.

Suppose that we have a tour traversing an equation gadget of the type shown in Fig. 10b in such a way that there are no semitraversed lambdas in it. If there are zero or two traversed lambdas, it is possible to modify this tour, without increasing its length and without changing the tour on the lambdas, in such a way that there are no endpoints in the gadget. Otherwise, it is possible to construct a tour with two endpoints in the gadget and impossible to construct a tour with less than two endpoints in the gadget.

Proof. Figures 11 and 12 show that there exist tours with the number of endpoints stated in the lemma. To complete the proof, we must show that it is impossible to construct better tours in the cases where the tour has two endpoints in the gadget.

If there is to be no endpoint in the gadget a number of conditions must be satisfied: The tour must traverse the edge to B, the edge from F, the edge DI, and the edge IM in Fig. 10. Thus we can assume that one partial tour enters the gadget through the vertex B and that another, or possibly the same, partial tour leaves through the vertex F. Since there are no semitraversed lambdas in the gadget, the only way, other than through the above described vertices, a partial tour can leave the gadget is through an endpoint, which in turn implies that there is an even number of endpoints in the gadget.

Let us suppose that the lambdas in the gadget are traversed as shown in Fig. 12a and that there exists a tour having no endpoints in the gadget. To avoid making the vertex M an endpoint, one of the edges KM and MO must be traversed. But since both lambdas are traversed, neither the edge KM or the edge MO can be traversed. This is a contradiction, which implies that it is impossible to construct a tour with zero endpoints. With a similar argument, we conclude that the same holds for the other cases shown in Fig. 12. □

Lemma 18. *Suppose that we have a tour traversing an equation gadget in such a way that there is exactly one semitraversed lambda in it. Then it is possible to modify this tour, without increasing its length and without changing the tour on the lambdas, in such a way that there is one endpoint in the gadget, and it is impossible to construct a tour with less than one endpoint in the gadget.*

Proof. From Fig. 13 we see that we can always construct tours such that there is one endpoint in the gadget. We now show that it is impossible to construct a tour with fewer endpoints. As in the proof of Lemma 17, we can assume that one partial tour enters the gadget at B and that another, or the same, leaves at F. Since there is one semitraversed lambda in the gadget, one partial tour enters or leaves the gadget at that lambda, which implies that there must be an odd number of endpoints in the gadget. □

Lemma 19. *Suppose that we have a tour traversing an equation gadget in such a way that there are two semitraversed lambdas in it. Then it is possible to modify this tour, without increasing its length and without changing the tour on the lambdas, in such a way that there are two endpoints in the gadget, and it is impossible to construct a tour with less than two endpoints in the gadget.*

Proof. From Fig. 14 we see that we can always construct tours such that there are two endpoints in the gadget. In order to prove the last part of the lemma we must argue that it is impossible to traverse the gadget in such

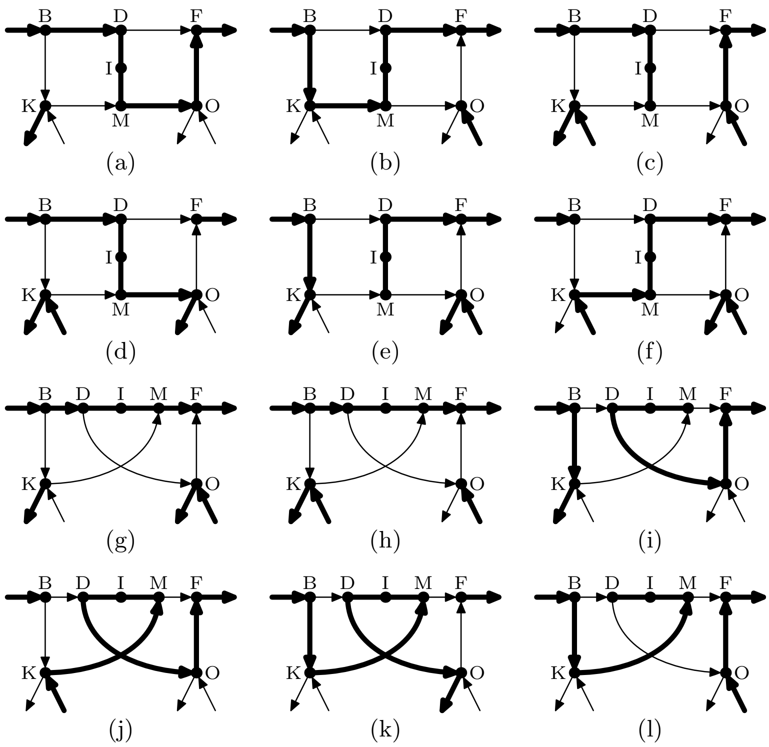

Figure 13. If one lambda in an equation gadget is semi-traversed, there must be at least one endpoint in it. In the cases where a lambda-vertex is also an endpoint, either one of the corresponding lambda-edges can be traversed by the tour. We show only a few of all possible combinations above.

a way that there are no endpoints in it. By an argument similar to that present in the proofs of Lemmas 17 and 18, a partial tour will enter (or leave) the gadget at four places, which implies that there is an even number of endpoints in the gadget. If there is to be no endpoints in the graph, there must be two partial tours in the gadget. Since the tours cannot cross each other and the gadget is planar we have two possible cases.

The first case is that the partial tour entering the gadget at B leaves it at F and the partial tour entering at K leaves at O. These two partial tours cannot, however, traverse the edges DI and IM without crossing or touching each other. As noted in the proof of Lemma 17, these edges must be traversed for the gadget to contain no endpoints. Thus, we can rule this case out.

The second case is that the partial tour entering the gadget at B leaves it at K and the partial tour entering at O leaves at F. Since these two partial tours cannot traverse the edges DI and IM, we conclude that at least two endpoints must occur within the equation gadget. □

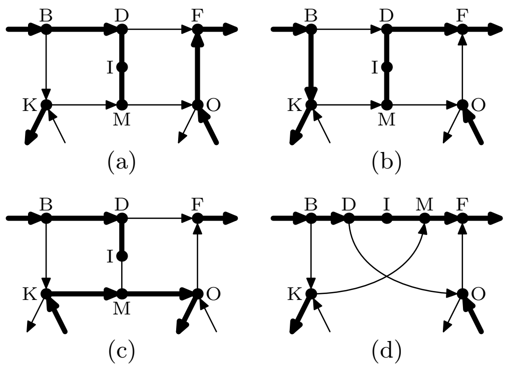

Figure 14. If both lambdas in an equation gadget are semi-traversed, there must be at least two endpoints in it. In the cases where a lambda-vertex is also an endpoint, either one of the corresponding lambda-edges can be traversed by the tour. We show only a few of all possible combinations above.

Lemma 20. *It is always possible to change a semitraversed lambda to either a traversed or an untraversed lambda without increasing the number of endpoints in the tour.*

Proof. First suppose that only one of the lambdas in the equation gadget is semitraversed. By Lemma 18 we can assume that the gadgets are traversed according to Fig. 13. Let us study the tour shown in Fig. 13a. By replacing it with the tour shown in Fig. 11a, we remove one endpoint from the equation gadget, but we may in that process introduce one endpoint somewhere else in the graph. In proof, let λ be the left lambda-vertex in Fig. 13a and v be the vertex adjacent to λ through the untraversed lambda edge. If v is an endpoint, we simply let the partial tour ending at v continue to λ, thereby saving one endpoint. If v is not an endpoint, we have to reroute the tour at v to λ. This introduces an endpoint at a neighbor of v, but that endpoint is set off against the endpoint removed from the equation gadget. To sum up, we have shown that it is possible to convert the tour in Fig. 13a to the one in Fig. 11a without increasing the total number of endpoints in the graph. In a similar way, we can convert the other tours shown in Fig. 13 to tours shown in Fig 11.

Finally, suppose that both lambdas are semitraversed. By Lemma 19 we can assume that the gadgets are traversed according to Fig. 14. By the method described in the previous paragraph we can convert the tours shown in Fig. 14 to suitable tours shown in Fig 13. As noted in the previous paragraph these tours can be converted to tours shown in Fig 11 at no extra cost. □

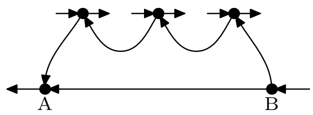

Figure 15. There is one lambda-vertex for each occurrence of each variable. The three lambda-vertices corresponding to one variable in the system of linear equations are joined together in a variable cluster. The three uppermost vertices in the figure are the lambda-vertices.

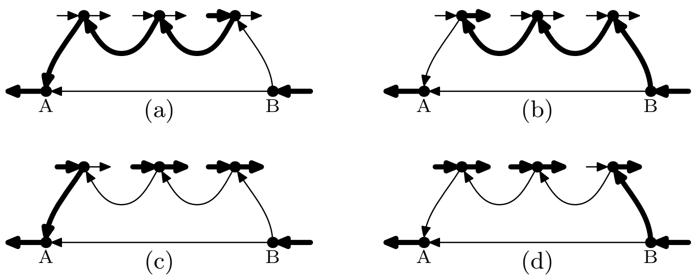

Figure 16. If the first or the last lambda in a variable cluster is semitraversed, the cluster is traversed as shown above.

4.2 The variable cluster

The asymmetric variable cluster is shown in Fig. 15. The edge from A in one cluster coincides with the edge to B in another, as described in Sec. 4.3. Suppose that the variable cluster corresponds to some variable x. Then the upper three vertices in the cluster are lambda-vertices in the equation gadgets corresponding to equations where x occurs.

Lemma 21. *Suppose that we have a tour traversing a cluster in such a way that there are some semitraversed lambdas in it. Then, it is possible to modify the tour, without making it longer, in such way that there are no semitraversed lambdas in the cluster.*

Proof. Suppose that there is one semitraversed lambda. If the semitraversed lambda is the middle lambda of the variable cluster it can, by Lemma 20, be transformed into either a traversed or an untraversed lambda. This moves the semitraversed lambda to the end of the cluster. Then the cluster looks as in Fig. 16. By moving the endpoint in the variable cluster to the

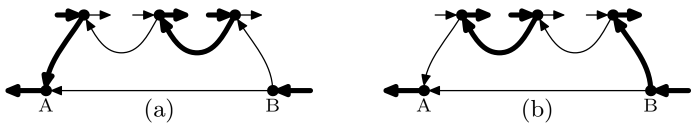

Figure 17. If all lambdas in a variable cluster are semitraversed, the cluster is traversed as shown above.

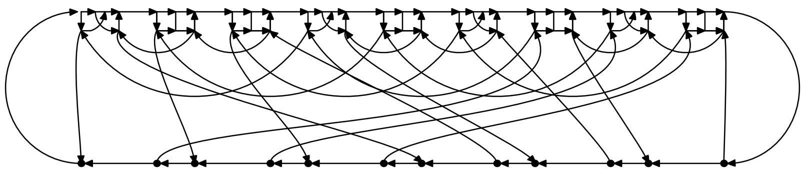

Figure 18. All asymmetric equation gadgets and all asymmetric variable clusters are linked together in a circular chain as shown schematically above. The equation gadgets are at the top of the figure and the variable clusters at the bottom. The precise order of the gadgets is not important. For clarity, we have omitted all vertices from the equation gadgets.

equation gadget corresponding to the semitraversed lambda, we can make the last semitraversed lambda traversed or untraversed without changing the number of endpoints.

Suppose now that there are two semitraversed lambdas. By Lemma 20, they can be transformed into either a traversed or an untraversed lambda without changing the number of endpoints in the tour. This implies that we can transform the tour in such a way that there is only one semitraversed lambda without changing the number of endpoints in the tour. Then we can use the method from the above paragraph to transform that tour in such a way that there are no semitraversed lambdas.

Finally, suppose that all three lambdas are semitraversed. Then the variable cluster would be traversed as in Fig. 17. By Lemma 20, the tour can be transformed in such a way that the two outer lambdas in the variable cluster are either traversed or untraversed without changing the weight of the tour. If the center lambda is not semitraversed after the transformation, the proof is complete. Otherwise we can apply the first paragraph of this proof. □

4.3 The entire Asymmetric (1,2)-TSP instance

The equation gadgets and variable clusters are hooked together in a circular chain as shown in Fig. 18. We complete our proof for the asymmetric case in the same way as in Sec. 3.3.

Lemma 22. *Given a tour with* $2e$ *endpoints, we can construct an assignment leaving at most* e *equations unsatisfied.*

Proof. Given a tour, we can by Lemmas 17–21 construct a new tour, without increasing its length, such that for each variable cluster either all or no lambda-edges are traversed. Then we can construct an assignment as follows: If the lambda-edges in a cluster are traversed by the tour, the corresponding variable is assigned the value one; otherwise it is assigned zero. By Lemma 17, this assignment has the property that there are two endpoints in the equation gadgets corresponding to unsatisfied equations.

Thus, the assignment leaves at most e equations unsatisfied if there are $2e$ endpoints. □

Theorem 23. *For every $\epsilon_1 > 0$ and $\epsilon_2 > 0$ it is **NP**-hard to decide whether an instance of the asymmetric traveling salesman problem with distances one and two with $2800n$ nodes has an optimum tour with length above $(2805-\epsilon_1)n$ or below $(2804+\epsilon_2)n$.*

Corollary 24. *It is, for any $\epsilon > 0$, **NP**-hard to approximate Asymmetric (1,2)-TSP within $2805/2804 - \epsilon$.*

Proof of Theorem 23. The result of Berman and Karpinski [1] states that it is **NP**-hard to determine if an instance of E2-Lin(3) mod 2 with $336n$ equations has its optimum above $(332-\epsilon_2)n$ or below $(331+\epsilon_1)n$. If we construct from an instance of E2-Lin(3) mod 2 an instance of Asymmetric (1,2)-TSP as described above, the graph in the Asymmetric (1,2)-TSP instance contains $25n$ nodes, given that the E2-Lin(3) mod 2 instance contains $2n$ variables and $3n$ equations. Thus, Lemma 22 and the above hardness result together imply that it is **NP**-hard to decide whether an instance of (1,2)-TSP with $2800n$ nodes has an optimum tour with length above $(2805 - \epsilon_1)n$ or below $(2804 + \epsilon_2)n$. □

5 Concluding remarks

We have shown in this paper that it is, for any $\epsilon > 0$, **NP**-hard to approximate Asymmetric (1,2)-TSP within $2805/2804 - \epsilon$ and (1,2)-TSP within $5381/5380 - \epsilon$. Since the best known upper bound on the approximability is 17/12 for Asymmetric (1,2)-TSP and 7/6 for (1,2)-TSP, there is certainly room for improvements. Our lower bound follows from a sequence of reductions, which makes it unlikely to be optimal. The sequence starts with E3-Lin mod 2, systems of linear equations mod 2 with exactly three variables in each equation. Then follows reductions to, in turn, E2-Lin mod 2, E2-Lin(3) mod 2, and Asymmetric (1,2)-TSP or (1,2)-TSP, respectively. Thus, our hardness result ultimately follows from Håstad's optimal lower bound on E3-Lin mod 2 [5]. Obvious ways to improve the lower bound is to improve the reductions used in each step, in particular our construction in this paper and the construction of Berman and Karpinski [1]. It is probably harder to improve the lower bound on E2-Lin mod 2, since the gadgets used in the reduction from E3-Lin mod 2 to E2-Lin mod 2 are optimal, in the sense that better gadgets do not exists for that particular reduction [5, 9]. Even better would be to obtain a direct proof of a lower bound on (1,2)-TSP. It would also be interesting to study the approximability of $\triangle$-TSP in general, and try to determine if $\triangle$-TSP is harder to approximate than (1,2)-TSP.

6 Acknowledgments

Viggo Kann contributed with fruitful discussions on this subject. Also, Gunnar Andersson and Viggo Kann gave valuable comments on early versions of the paper. One of the anonymous referees spotted a lapse in a previous version of this paper and helped improve the presentation of the results. Santosh Vempala suggested that I should study also the asymmetric version of (1,2)-TSP.

References

1. Piotr Berman and Marek Karpinski. On some tighter inapproximability results. Technical Report TR98-029, Electronic Colloquium on Computational Complexity, June 1998.
2. Nicos Christofides. Worst-case analysis of a new heuristic for the traveling salesman problem. Technical Report CS-93-13, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, 1976.
3. Wenceslas Fernandez de la Vega and Marek Karpinski. On approximation hardness of dense TSP and other path problems. Technical Report TR98-024, Electronic Colloquium on Computational Complexity, April 1998.
4. Dimitris A. Fotakis and Paul G. Spirakis. Graph properties that facilitate travelling. Technical Report TR98-031, Electronic Colloquium on Computational Complexity, June 1998.
5. Johan Håstad. Some optimal inapproximability results. In *Proceedings Twenty-ninth Annual ACM Symposium on Theory of Computing*, pages 1–10. ACM, New York, 1997.
6. David S. Johnson and Christos H. Papadimitriou. Computational complexity. In Eugene L. Lawler, Jan K. Lenstra, Alexander H. G. Rinnoy Kan, and David B. Shmoys, editors, *The Traveling Salesman Problem*, chapter 3, pages 37–85. John Wiley & Sons, New York, 1985.
7. Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, New York, 1972.
8. Christos H. Papadimitriou and Mihalis Yannakakis. The traveling salesman problem with distances one and two. *Mathematics of Operations Research*, 18(1):1–11, February 1993.
9. Luca Trevisan, Gregory B. Sorkin, Madhu Sudan, and David P. Williamson. Gadgets, approximation, and linear programming. In *Proceedings of 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 617–626. IEEE Computer Society, Los Alamitos, 1996.
10. Sundar Vishwanathan. An approximation algorithm for the asymmetric travelling salesman problem with distances one and two. *Information Processing Letters*, 44:297–302, 1992.